

A MISSION OPERATIONS ARCHITECTURE FOR THE 21ST CENTURY

W. Tai and D. Sweetnam

*Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, California 911090
818-354-7771 (v), 818-354-9068 (f)
wtai@jpl.nasa.gov, dsweetnam@jpl.nasa.gov*

ABSTRACT. A new operations architecture for low cost missions is proposed, to build a framework to enable NASA to face its major challenges beyond 2000. This architecture is composed of three elements:

- Service Based Architecture ● Demand Access Automata •Distributed Science Hubs

The Service Based Architecture is predicated on a set of standard multimission services being defined, packaged and formalized by JPL's Deep Space Network and Advanced MultiMission operations System. It is analogous to the services paradigm of a telephone company, Demand Access Automata is a suite of technologies that break today's contact conundrum, the cost driving requirement for nearly continuous contact with each and every craft we fly. We describe a spacecraft initiated 'beacon', a 'virtual emergency room', and a 'high efficiency tracking'. Distributed Science Hubs provide a trio of information system capabilities to the small science oriented flight teams of the future, consisting of individual access to all traditional mission functions and services, powerful multimedia intrateam communications to facilitate collaborative investigation, and automated direct transparent communication between scientist and instrument.

1. INTRODUCTION

Leading into the 21st century, the planetary program for NASA faces two primary challenges, first to increase the rate of exploration missions, and second to live with budget allocations that may be constant or even declining. To respond to these challenges, NASA has taken the initiative on several fronts. First, NASA has begun a low cost exploration program, Discovery, to motivate a larger number of missions by keeping them small, with focused science, low cost and increased scientist participation. Second, NASA has begun an advanced technology development wedge to provide more capable devices at less cost. Third, NASA is endeavoring to reduce the cost of operating the missions it flies, by a combination of more efficient processes and automation.

It is the purpose of this paper to describe an architecture to meet these challenges. We have established objectives as follows:

- . Reduce the cost of Mission Operations System (MOS) development by a factor of 10.
- . Reduce the cost of Operation by a factor of 10.
- . Reduce the time of MOS development by a factor of 10.

Our approach to meeting these objectives is to establish an end-to-end mission operations architecture based on the adoption of standards, the use of efficient processes, and the incorporation of advanced technology. We have three major components, first the provision of standard mission operations services that can be contracted for by Project customers, second the development of technologies to automate key communications interfaces, and third an implementation of an underlying standard information system infrastructure that makes the investigator to instrument link transparent. Figure 1 provides an overview of this architecture.

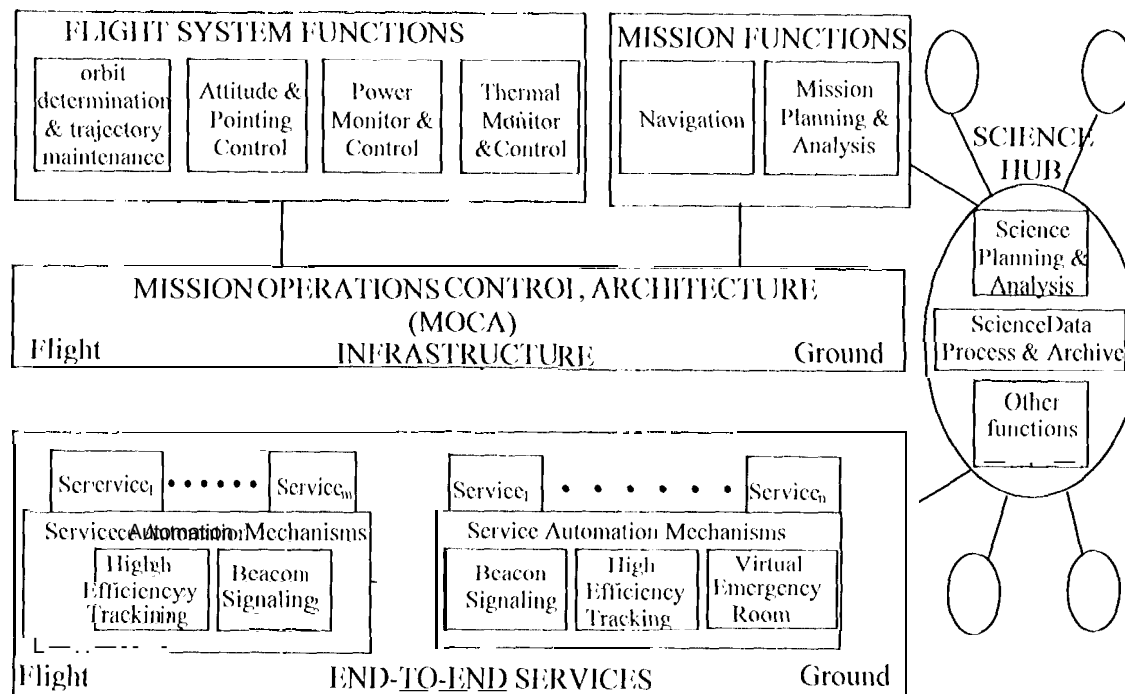


Figure 1. End-to-End Mission Operations Architecture

2. SERVICE ARCHITECTURE

BACKGROUND. JPL deep space missions have traditionally used some mission operations services provided by the two JPL multi-mission systems, i.e. the Deep Space Network (DSN) and the Advanced Multi-Mission Operations System (AMMOS), both under the management of the JPL Telecommunications and Mission Operations Directorate (TMO). The AMMOS also supplies multi-mission tools to flight projects to build project-unique Ground Data Systems (GDS). It has been observed that different flight missions utilize multi-mission services and tools in varying mix ratios. Furthermore, the ratio of the subscribed multi-mission services to the project-developed MOS capabilities differs widely from mission to mission. As NASA moves into 21st century, its full cost accounting policy, which will be in place in the near future, has some key implications to both flight missions and multi-mission systems. To flight missions, duplication of effort between project MOS and multi-mission systems must be avoided. The most cost-effective approach to developing MOS is to put their precious dollars onto truly mission-

specific capabilities and, in areas where both multi-mission choices (i.e. services and tools) are available, subscribe to multi-mission services as much as possible. To the multi-mission systems, while there will be an increase in demand for their services, full cost accounting means services provided must be as cheap and responsive as possible, otherwise services will no longer be justifiable. The bottomline is that multi-mission systems have to move into a more service-oriented paradigm than they are today.

WHY STANDARD SERVICES? To meet the development duration and cost objectives as discussed in the Introduction section, a service architecture must be in place to accommodate the end-to-end architecture of the 21st century. A key concept of the architecture is the standard services provided by the DSN and AMMOS. These standard services are formalized for the following purposes: (1) as "instant" services to support flight projects from the start of the development phase so that they can focus their effort on mission-specific capabilities, (2) reduce mission risk liability by providing service-domain expertise/talents since projects can no longer afford to keep service-domain expertise/talents like they do today, (3) standard services, each defined and packaged in a Service Catalog for performance and cost accountability, like a phone company 'Calling, Services' menu, will make full cost accounting easier (4) individual multi-mission services formalized as standard services will be selectable by each mission for its operational needs, (5) formalization of standard services to be cost competitive will motivate the DSN and AMMOS to automate its operations.

DEFINITION AND KEY ATTRIBUTES OF STANDARD SERVICES. We have defined the term "service" in a rather limited way for the purpose of characterizing the service paradigm. A service is work performed by the multi-mission systems, using one or more tools, facilities, (and staff), that produces mission operations results for a customer (i.e. flight project or science investigator). Services may be "standard" or "tailored". Standard services are those defined in the Service Catalog from which customers can make selection for their needed operations to support their missions without the expenditure of non-recurrent engineering. A tailored service is one requested by customers for functionality different from a corresponding standard service offered in the Service Catalog and, for fulfilling this service, modification of existing capabilities with additional implementation effort will be needed. In the service paradigm, the standard services have the following key attributes:

- (1) **RELEVANCE:** Services offered to the customers must be visible and meaningful to the customers. This implies hiding the level of details of the capabilities and activities from the customers.
- (2) **PICK-AND-CHOOSE:** The services must be selectable by customers. Subscription to a service by a customer should not require buy-in of other services which are not relevant to the customer's needs.
- (3) **PLUG-AND-PLAY:** The use of any standard services (as distinguished from the tailored services) must be based on definitions which appear in the Services Catalog. Once a service, as it exists on the Services Catalog, is subscribed to, it must be "immediately" available for use by the customer. It should not require any implementation effort beyond interface testing, configuration setup, and parameter table updates, by the services provider.
- (4) **STANDARD INTERFACES:** The use of the standard services, in terms of control and data interfaces, by the customers will be via standard interfaces. "Standard" interfaces include those formally established by standards organizations, those widely applied by the industry as de facto standards, and those defined by the service provider as common

mechanism to all customers. No additional development effort on the multi-mission or the subscriber's system other than that required for conforming to the standard interfaces will be necessary.

- (5) **DIRECT SERVICE CONTROL:** The customers will be allowed to directly control the service (within the bounds of the system's capabilities).
- (6) **INTEROPERABILITY:** Services will be standardized, whenever applicable, to enable interoperability with other service providers whenever the same service is requested.
- (7) **PERFORMANCE ACCOUNTABILITY:** Performance of each individual service subscribed to by a customer will be measurable and reportable.
- (8) **COST ACCOUNTABILITY:** Services will be provided to a customer on a fee schedule basis. This means all standard services will be defined, structured, and priced in such a way that customers' recurrent costs can be tracked and reported to them.

DESCRIPTION OF SERVICE SYSTEM ARCHITECTURE. Figure 2 is a diagram describing the service architecture in a layered view of the end-to-end system. At the highest layer, i.e. the function layer of project mission operations and flight system, there are a set of mission unique functions. These functions, flight- and ground-based, play together utilizing the underlying end-to-end standard services at the next lower layer to accomplish the project's mission objectives. The standard services layer is further divided into 2 sub-layers: application service sub-layer and data acquisition/delivery/management service sub-layer. Services in the former are more mission-oriented applications whereas those of the latter sub-layer deal with the acquisition, delivery, and management of mission data without any concern of the mission-dependent data contents. 12 families of standard services have been defined and documented in the Service Catalog. Each service family contains one or more individual services starting from lower level products to higher level. These individual services in the service families constitute the building blocks of the service architecture. For example, the telemetry service family includes frame service, packet service, channel processing service, and data set service. Customers can subscribe to frame service without having to subscribe to the other services. And the selection of higher layer services, e.g. packet service, does not require the explicit selection (or additional service fee) of frame service since the frame-level processing is inherently part of the functions performed for packet service. The tracking and navigation service family also offers from the lowest level to higher level of services: radiometric service, orbit determination service, and trajectory analysis service.

A special family of services is called service management services which are defined based on the Cross Support Reference Model for CCSDS Space Link Extension (SLX) Services. These are the services conducted by the multi-mission systems to facilitate customer's request for services of any other 11 families. They include allocation and scheduling of resources, configuration and control of assets, service accountability reporting (for performance and cost), etc..

From the end-to-end system perspective, 4 of the service families will be offered in the near future as end-to-end services: command services, telemetry services, tracking and navigation services, and mission data management services, although the provision of the others can also be migratable between ground and flight. The concept of end-to-end services is based on the need to avoid spacecraft development costs wherever possible and the feasibility that the on-board functions such as telemetry frame encoding (Reed-Solomon and Turbo code), command frame decoding, CCSDS COP-1, and timing data in addition to the traditional deep space transponder, and their counterparts on ground, can be "pre-fabricated", integrated and tested. Another

beneficial result of the end-to-end services is they enable the implementation of various demand across and high-efficiency tracking schemes, thus accommodating the kind of autonomy requiring non-deterministic orchestration between flight and ground.

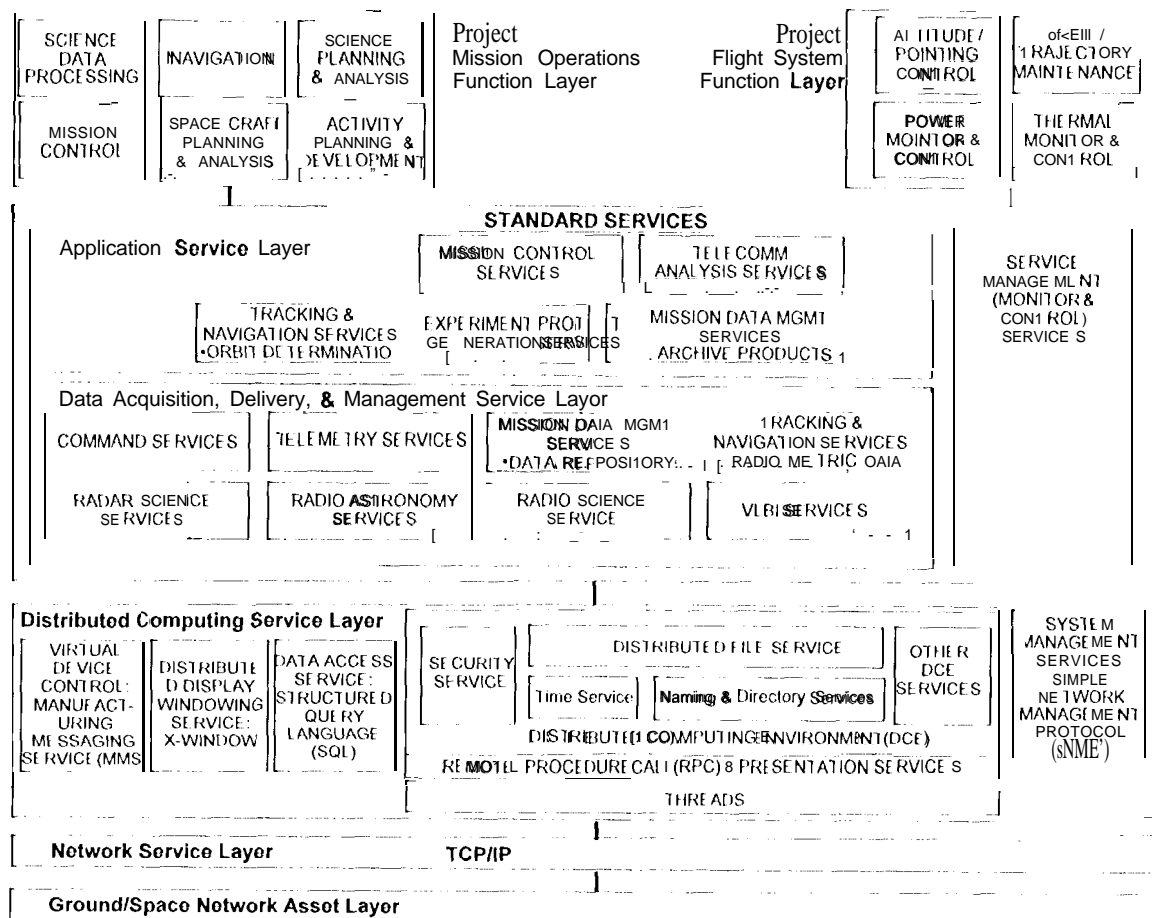


Figure 2. Service System Architecture - A layered View

3. DITMANJ ACCESS AUTOMATA

BACKGROUND. In this section we describe a suite of technologies collectively known as demand access automata, designed to dramatically reduce the cost of contact with spacecraft. The components are a spacecraft 'beacon' to initiate contact only when necessary, a ground 'virtual emergency room' to orchestrate the ground response, and a 'high efficiency tracking' scheme to optimize the telemetry link itself. They provide for complete automation of spacecraft - ground communications and radically alter traditional processes by establishing the spacecraft as the initiator of communication services.

Analysis of the tracking patterns has shown that an enormous amount of time and money has been spent in contact (routine or otherwise) with NASA planetary spacecraft. As examples, the Mars observer spacecraft, on its cruise to Mars in 1993, was in contact with the ground for an

average of 22 hours per day; the Magellan spacecraft, in orbit around Venus in 1993, was in contact an average of 26 hours per day (use of 2 antennas raises the level above 24 hours). When combined with tracking cost rates from the Deep Space Network for 34 meter antennas, the estimated costs for NASA to track these spacecraft are \$9.9 million dollars per year and \$11.8 million dollars per year, respectively.

BEACON SIGNALING. 'Beacon' signaling is a low cost mechanism to allow spacecraft to initiate contact with the ground. See figure 3. The spacecraft, using an intelligent onboard health/safety monitoring agent, assesses its own health. It then filters that assessment into one of 3 'beacon' states, essentially, "I'm OK", "I need 1111,1", or "I want to DUMP telemetry". Using a simple RF signaling scheme, such as a pair of frequency subcarriers with a high modulation index to suppress the carrier, and with the frequency difference being related to the desired 'beacon' state, the signal is transmitted to Earth. Since little power is required to transmit such small amounts of information, most spacecraft can use broad beam low gain antennas to radiate the 'beacon' signal while still focusing on scientific observations or cruising.

The 'beacon' signals from any number of equipped spacecraft are monitored by a small network of dedicated 'beacon' reception stations. These are small (perhaps 3-11 meter) low cost receive-only antennas that search the sky once each day polling each spacecraft for its state. They are equipped with adaptive weak-signal detectors that search, in turn, for one of the 3 'beacon' states. Although one must search in location, frequency, and frequency rate, after making provision for the motion of transmitter and receiver, this task is simplified by the small number (3, as noted above) of known spacecraft states to look for. The 'beacon' reception station provides an output for each of the spacecraft that it has been asked to look for. This is a probability of detection of each of the states, with a very high probability indicating successful detection; this means that a fourth state, "NO detection" can also be forwarded to the 'Virtual Emergency Room'.

VIRTUAL EMERGENCY ROOM. The 'Virtual Emergency Room' (VER) is the central facility for directing the activities of the 'beacon' reception stations and orchestrating the appropriate ground response to all of the possible 'beacon' states for a particular spacecraft. It is fully automated. Information about every spacecraft to be polled is maintained in a database. This includes locations and frequency tuning profiles. Information is updated daily after successful contact has been made. This database would also contain the duty roster of people to contact in case of an emergency.

In addition, a second rule-base contains the specific responses for each spacecraft for each possible 'beacon' state. An example rule set might be to send an e-mail message to the Project Manager after detecting the "I'm OK" state and take no further action. In the event that the "1111,1" state is detected, the VER would notify key people on the duty roster with a paging system and then initiate a request for an emergency tracking pass using a large (34-70 meter) antenna to obtain an engineering status packet. In the event that the 'DUMP' state has been activated, the VER would initiate a request for a 'High Efficiency Tracking' pass to damp telemetry. The Project (if the spacecraft is not in a position to communicate with Earth) could set a rule that 'NO detection' for up to 'n' days is the same as "I'm OK".




Figure 3. Service Automation Mechanisms - Spacecraft Initiated Access

HIGH EFFICIENCY TRACKING. The 'High Efficiency Tracking' technology provides the capability for a spacecraft to downlink telemetry at the best data rate, determined in near real-time at the beginning of the telemetry pass, not months or years in advance as is current practice. When the VÉR initiates a request for a 'High Efficiency Tracking' pass (see above), a large antenna (34 or 70 meter) resource is allocated as soon as available. The antenna then assesses its own current condition, including communication frequency, tracking path, weather, and system noise temperature. It then prepares a message containing the best estimate of its performance capabilities during the next 8-12 hours, transmits this to the spacecraft and waits a round trip light time.

On receipt of the message, the spacecraft assesses its own performance, including space loss, transmitter power, and pointing losses. It combines this with the ground antenna performance message and constructs a variable bit rate profile for the next 8-12 hours. The spacecraft proceeds to send telemetry according to this rate profile. The ground antenna will record the full telemetry spectrum and then perform optimal bit detection using iterative digital receiver and detector algorithms (note that this can introduce a delay in the telemetry stream of seconds to minutes). The telemetry stream is then forwarded to the Project. At the end of the transmission, the spacecraft returns to the 'beacon' signal, "I'm OK".

We make mention of a Mission Operations Control Architecture (MOCA). MOCA is a specification of mechanisms required to conduct a telecommunicated process control dialog between a human (or machine) user and remote, distributed space mission systems. Its major

components are 1) the Control Interface which allows the human (or automated) controller to specify and monitor the desired sequence of operations to be conducted in a remote system, 2) the Space Messaging System which translates machine-readable command calls from the Control Interface into standard-syntax messages that invoke the desired actions in the remote end system and return standard-syntax response information to the controller, and 3) the Decision Support Logic which allows rules for command execution to be programmed into a distributed inference engine (which may be located wholly on the ground, wholly in space, or partitioned in varying ways between the two).

MOCA provides an infrastructure that has two key benefits for our End-to-End architecture. First, through its Space Communications Protocol Standards (SCPS) capabilities, there is provision for an investigator to directly access his instrument. In simplest terms, this is provision of an FTP capability directly to and from the instrument originating and terminating in the Science Hub at the scientist's workstation (taking into account possibly long light times). Second, the MOCA Control Interface with SCPS provides a common protocol and standard for simultaneous control of both the spacecraft subsystems and the ground tracking subsystems. This is a necessary building block for the Service Automation mechanisms, such as Beacon Signaling, High Efficiency Tracking.

4. SCIENCE HUBS

Future small missions, such as those in the Discovery Program, will be characterized by geographically distributed flight teams of 10 (or less) scientists and engineers. This is a departure from previous missions which have stressed co-location of the flight team that may reach as large as several hundred engineers. In addition, there has been a tendency to isolate the science community from the actual operation of the mission, relegating it to high level requests for observation time and squabbles over data formats, in addition to the detailed analysis of data and the publication of results.

In our architecture, we realize that the cost of co-location of operations team members is going to be prohibitive. We also realize that NASA, in restructuring its programs toward smaller focused science objectives, is demanding higher levels of participation by scientists. We are also aware of the information system technology explosion in areas such as continental and intercontinental communication and multimedia.

A Science Hub represents a logical grouping of scientists (and engineers) to conduct a specific mission (see figure 1). As NASA is able to increase the launch rate for new missions, the number of science hubs grows. The Hubs represent the focus of activity for this particular mission, with participants further distributed around the Hub. A simple view would be that the Principal Investigator represents the Hub, with (h) investigators and engineers distributed around the Hub.

The Science Hub gains access to the spacecraft and instruments by contracting for the 'End-to-End Services' (see section 2). Of course, choice of the specific services is selectable, according to the mission needs. This leads to an extension of the MOCA infrastructure to the science hub by means of network connections, multimedia workstations and software. The software will take the form of a "Client Package", which provides all of the selected interfaces to the service provider "Servers" and access to the selected service functions. A key benefit of this approach is

that this “MOS development” is accomplished using “off the shelf” services enabling installation in a few months instead of a few years.

A functional view of how the MOS would appear to the scientists and engineers in the Science 1 hub is shown in Figure 4. In addition to the “Client Package”, there is provision for mission functions (such as the “Science Planning and Analysis” or “Science Data Processing & Archiving”). There is provision for additional “mission processes” (which may be aggregations of services and custom science analysis packages). Of course, there is access to a comprehensive suite of office tools (word processors, spreadsheets, drawing/presentation, visualization), and to Internet browsers.

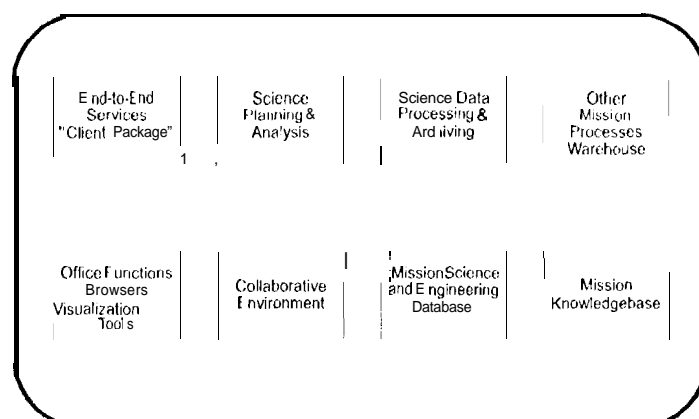


Figure 4. Functional view of Science Hub user interface

We provide a “collaborative environment” which is the mechanism for efficient communication between the geographically distributed team of scientists and engineers. This contains real-time visual and audio, as well as distributed whiteboards and application displays. Finally, we provide for a “knowledge base”, which is the repository of the mission requirements, design, drawings, testing, operational procedures, and lessons learned during the course of the mission. It would contain documents as well as depositions from the designers, spacecraft builders and flight operators about how the mission was put together.

S. CONCLUSIONS

IMPLICATIONS. We can reduce our development time by a factor of 10 by going from a typical GDS development time of 30 months to something like 3-4 months. We would expect that the cost of GDS implementation for a specific mission will drop to 1-3 million dollars, depending on selected services. We do this by providing new missions with selections of off-the-shelf multi-mission services that require no development...only selection, integration, and testing.

We can reduce the cost of operations a factor of 10 by using flight teams of 10 people or less, composed of scientists, mission engineers, and spacecraft engineers, using the standard services from their science hub. This small flight team size is achievable, in part, due to NASA's emphasis on small focused science programs like Discovery. A major contribution, however, to reduced team sizes, will be the result of the service automation mechanisms. We estimate that tracking hours with large expensive antennas can be reduced by 90% during cruise periods by use of 'Beacon' signaling, and that tracking coverage for telemetry downlinking can be reduced by at least 50% by using "High Efficiency Tracking".

PROGRESS. The Jet Propulsion Laboratory is currently developing this End-to-End architecture and is prototyping the technology components we have described, using testbeds provided by '1 MOJ). In addition, the New Millennium DS 1 mission is expecting to demonstrate a version of the 'Beacon' signaling concept. It is our expectation that NASA will implement full cost accounting and Total Mission Cost (which will require estimates of the cost of using any or all of the DSN and AMMOS services) in the next couple of years, followed by direct charging to Projects of the full cost of these services. This push will certainly provide the motivation to complete development of this architecture, and its automation technologies and propel us into the next century.

The research described in this paper was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.